Master-Seminar - Hochleistungsrechner: Aktuelle Trends und Entwicklungen (IN2107, IN4835)
Winter Term 2015/2016

Use of mobile SoCs in HPC

Paul Nieleck
Technische Universität München

3.2.2016

Abstract

High Performance Computing (HPC) systems consume a lot of power. It is assumed that in some years the first systems that have a performance of 1 EFLOPS[1] will appear. Due to financial and physical restrictions such a system shall consume not more than 20 MW of electrical power. Therefore it is necessary to massively increase the energy-efficiency of supercomputers. Using low-energy mobile processors for supercomputers to reach this goal is currently under active research. This paper presents results from different scientific publications regarding this issue. Comparing the energy-efficiency of ARM Cortex-A8 and Cortex-A9 CPUs from single-board computers with Intel Xeon E5 CPUs in terms of MFLOPS/W shows that these ARM CPUs are less energy-efficient. The Cortex-A9 CPU of the ST-Ericsson Nova A9500 System-on-Chip (SoC) has an energy-efficiency of 235.0 MFLOPS/W. An Intel Xeon E5-2679 designed for high performance computing has an energy-efficiency of 674.69 MFLOPS/W. Nevertheless, there are already ARM-based server systems available on the market. The Viridis server system featuring many ARM cores has a similar energy-efficiency as server systems with Intel processors, but it has a more than ten times lower performance. Also some ARM-based supercomputers are currently being built. A first prototype is the European Mont-Blanc project. This paper gives an overview of the announced ARM-based HPC systems.

[1] 10^{18} floating point operations per second

1 Introduction

Energy consumption is an important factor in High Performance Computing (HPC). With increasing computing power the energy consumption often increases too. For example, the number one supercomputer on the *TOP500* list from November 2015 [19], the Tianhe-2 of the National Supercomputer Center in Guangzhou, China, consumes a power of 17,808 kW at a peak performance of 54,902.4 TFLOPS. Five years earlier, the number one supercomputer on the TOP500 list from November 2010 [18], the Tianhe-1A of the National Supercomputing Center in Tianjin, China, consumed a power of 4,040 kW at a peak performance of 4,701.0 TFLOPS. Although the ratio between performance and power is increasing more and more, with the goal of breaking the 1 EFLOPS barrier in mind, the energy consumption eventually will reach an physically and financially unacceptable figure. In their 2008 study about exascale computing in the year 2015, Kogge et al. [24] assumed a maximum of

20 MW electrical power consumption for the computational part of a supercomputing system. For comparison, typical onshore wind turbines with a rotor diameter of about 100 meters can produce about 3 MW of electrical power [14]. That means it would be necessary to operate seven turbines to power one supercomputer. Therefore it is necessary to find new solutions to increase energy-efficiency and counter the problem of massive energy consumption.

One such solution is to use accelerator hardware that is designed for special purpose computation, e.g. floating point operations, but is very energy-efficient. The success of this approach can be seen in *The Green500 List* [17], which appears twice a year and lists the most energy-efficient supercomputers in the world. In the November 2015 list, the 40 most energy-efficient supercomputers all use accelerator cards like NVIDIA Tesla GPGPUs or Intel's Xeon Phi coprocessors. By using accelerator hardware, the number one most energy-efficient supercomputer, "Shoubu" of the Institute of Physical and Chemical Research (RIKEN), Japan, reaches an energy-efficiency of 7031,58 MFLOPS/W. Nevertheless, this figure is still far away from the 50 GFLOPS/W that would be necessary to run a 1 EFLOPS supercomputer with a power consumption of 20 MW.

Another solution currently under active research is the use of mobile processors. Instead of CPUs designed for servers and supercomputers, processors like ARM or Intel Atom are used to build HPC clusters. These processors usually feature a low energy consumption by design, which makes them an object of research to target the energy consumption of supercomputers. Furthermore these processors are widely available at a considerably low price. Since they are consumer products that are used in tens of millions of devices, economy of scale has a positive effect on the production costs.

A huge downside of many mobile processors is that they are only 32-bit processors, limiting the memory size to 4 GB. Recently this is beginning to change as ARM introduced its 64-bit ARMv8-A architecture.

In the following sections focus is set on ARM processors, since they are the most common processors in mobile devices nowadays and most research is based on them. Section 2 will give an overview of current mobile Systems on Chips (SoCs) that feature mobile processors. In Section 3 the performance to energy ratio of several mobile processors is compared to the ones of common server processors that are currently used for HPC. Section 4 gives an overview of current HPC and server systems that are based on ARM processors. The paper is concluded with Section 5.

2 Mobile SoCs Overview

The wide acceptance of mobile platforms like smartphones, tablet PCs and laptops in the consumer market opened a huge market for companies developing processors for such platforms. Nevertheless there are essentially only two companies currently left competing in this market. One company is ARM Ltd., the other one is the Intel Corporation. Processors from ARM use a RISC architecture, while Intel processors use a CISC architecture. Blem et al. [22] have compared ARM and Intel instuction set architectures (ISA) and have shown that the used ISA is not important when it comes to energy efficiency.

2.1 ARM

ARM does not produce processors on its own. Instead, it only owns the intellectual properties of its processor designs, which it licenses to manufactures of mobile SoCs like Samsung, Texas Instruments, AMD and others. The portfolio of ARM consists of different architectures that are aimed at different use cases. For compute intensive tasks two instruction set architectures are available: the 32-bit ISA ARMv7-A and the 64-bit ISA ARMv8-A. Especially the ARMv8-A-architecture, supporting the 64-bit ISA and a larger set of registers, is suited for the use in servers and supercomputers. All of ARM's Cortex-A Series processors are of either the aforementioned ISAs.

Many of the Cortex-A processors have floating

point units (VFPvX architecture) and support SIMD (single instruction, multiple data) instructions, called NEON. On some of the Cortex-A processors these are optional, e.g. on the Cortex-A9. The floating point units add 16 double precision registers on Cortex-A8 [3], Cortex-A9 [4], or 32 double precision registers on Cortex-A15 [2].
Besides CPUs ARM also provides GPUs, called Mali. They are, like the CPUs, designed for low energy consumption and are available on many SoCs that feature ARM-CPUs. Some of the Mali-GPUs support OpenCL, which enables them to be used as GPGPUs. There are also mobile GPUs from other vendors which support OpenCL. Using GPGPUs is especially helpful in speeding up floating point operations. Many of the TOP500 supercomputers use GPGPU accelerator cards from NVIDIA (Tesla K20/K40, etc.). This shows the importance of such accelerators. Having such accelerators already available in the SoC has the advantage over external accelerator cards, that they are physically close to the CPU and can be connected directly. External devices need to be connected via the PCI-bus or similar buses. The direct connection reduces the latency and increases the throughput. Furthermore the GPUs of the mobile SoCs usually share the same memory with the CPUs, instead of having their own memory. Thus there's no additional data transfer from the main memory to the graphics memory necessary.

2.2 Intel

Intel tried to enter the market for low energy CPUs with its Atom architecture. Atom CPUs are used in low-end desktop PCs and laptops but didn't manage to gain a greater market share on smartphones and tablets. Nevertheless Intel also released Atom-CPUs for the server market. All Atom CPUs support the x86 32-bit ISA. Besides that, the server CPUs also support the x86_64 64-bit ISA. The usage of the long established x86 instructions has two advantages over the ARM instruction set:

1. Programs written for Intel compatible CPUs don't need to be rewritten/recompiled to run on this architecture. Hence, programs and libraries can be easily reused.

2. Compiler designers have a long gained experience in enabling their compilers of automatically optimizing code for the x86 architectures. Code that is highly optimized is also more energy-efficient, since it needs fewer instructions for the computational tasks. Compilers for the ARM architecture currently seem to be less advanced in automatic code optimization. For example for the benchmarks done in [27] the GCC (GNU Compiler Collection) compiler flags for ARM specific optimization showed little to no effect on the performance.

3 Performance vs. Energy

As a first step in order to find out whether mobile SoCs are a good choice for HPC, the performance of single cores/CPUs is compared to the performance of common server CPUs. This of course can not be done in terms of frequency or floating point operations per second (FLOPS) since the mobile SoCs usually have far less total computing power than server CPUs. Therefore the FLOPS per watt shall be compared. This should give a first hint on the capabilities of mobile SoCs.
The reason why FLOPS/W are used as performance measure, is because most applications in HPC are simulations in physics, geo science, engineering or finance, which rely on floating point operations.[2] For that reason this measure is also used for the Green500 list.
Although the aim of this paper is to figure out if mobile SoCs are suitable for building energy-efficient exascale supercomputers in terms of FLOPS, there might be applications for which floating point operations are less important. For example, with the growing importance of big data analysis, graph

[2]For example the supercomputer "SuperMUC" at the Leibniz Supercomputing Center, Garching, Germany, is mainly used to study "astrophysics and plasma physics, earth and environmental sciences, life and material sciences, engineering and computational fluid dynamics, and high energy physics" ([21]).

traversal finds its way into HPC. Such systems have different requirements and the question if ARM processors are more energy-efficient for such applications is therefore out of scope of this paper. The *Graph 500* list [15] and the *Green Graph 500* list [16] list the fastest/energy-efficient supercomputers available for graph traversal.

3.1 ARM Cortex-A8

Padoin et al. [27] tested the ARM Cortex-A8 CPU on a BeagleBoard featuring the OMAP3530 from Texas Instruments. The OMAP3530 has a single core ARM Cortex-A8 CPU running at a frequency of 600 MHz. For the measurement of the energy consumption they used a power analyzer connected to the BeagleBoard. Measurement was done with serial High-Performance Linpack (HPL). They achieved a performance of 19.749 MFLOPS/W. Unfortunately this value is not very useful, since the energy consumption of the whole system, instead of the single CPU, was measured.[3] Thus, the measurement contains the power consumption of components that might not be used on supercomputer nodes or should be measured on their own, e.g. the 512 MB DDR3 RAM, the on-board flash storage or the USB and HDMI connectors. This makes it hard to compare to measurements of other CPUs.

3.2 ARM Cortex-A9

In [27] Padoin et al. additionally did research on the ARM Cortex-A9 CPU. The tests were performed on a cluster of eight PandaBoards, each featuring an TI OMAP 4430 dual-core ARM Cortex-A9 CPU at a frequency of 1 GHz. Measurement was done with parallel High-Performance Linpack (HPL) on a 1000×1000 matrix. In their setup they achieved a performance of 92.041 MFLOPS/W. But, like with the BeagleBoard, they measured the energy consumption of a whole PandaBoard, instead of just measuring the energy consumption of the CPU. Thus, this measurement also contains an overhead generated by the RAM, connectors, card-readers, etc.

A later work of Padoin et al. [26] also examined, besides the PandaBoard, the Snowball single board computer by CALAO Systems SAS [6] and the Qseven compliant boards of the Tibidabo system [29] from the Mont-Blanc project. The Snowball features a ST-Ericsson Nova A9500 dual-core CPU @ 1 GHz and the Tibidabo boards feature a NVIDIA Tegra 2 dual-core CPU @ 1 GHz. With HPL on a 5000×5000 matrix, they achieved the following results for their measurements: The PandaBoard has a performance of 91.3 MFLOPS/W. This is almost the same performance that they achieved during their measurements in [27]. The slight difference might be related to the different input size of the test. The Snowball board achieved a performance of 235.0 MFLOPS/W and the board of the Tibidabo system achieved a performance of 161.5 MFLOPS/W. These figures show that the performance values can hugely differ even for the same processor types (Cortex-A9), depending on the overall system. Comparing Cortex-A9 processors based on measurements of whole system's energy consumption can only give a hint on the energy-efficiency of the processors.

Cortex-A9 CPUs can consist of up to four cores. In an early phase of the Mont-Blanc project (see Section 4.2) Rajovic et al. [30] analyzed the NVIDIA Tegra 2 and Tegra 3 platforms with different applications. Among them were applications for *Vector Operation*, *Dense matrix-matrix Multiplication* and *Fast Fourier Transform*. The Tegra 3 platform consists of a quad-core Cortex-A9 running at 1.3 GHz and therefore has a much better performance than the Tegra 2 platform with its 1 GHz dual-core CPU. The maximum power consumed to solve the given problems is almost the same for both platforms. But, due to the quad-core configuration, the time needed to solve the problem was much faster. On average the energy-to-solution of the Tegra 3 was only 67% of the energy-to-solution of the Tegra 2. Rajovic et al. argue that a higher density of cores in the ARM CPUs leads to a better energy-efficiency.

[3] This fact is not stated explicitly, but in [25], [28], and [26] Padoin et al. described that they measured the instantaneous power of the whole board.

3.3 ARMv8 processors

Processors that are based on the ARMv8 architecture just entered the market. For that reason, only a few research papers on energy-efficiency of these processors are available. None of these papers provide values for energy-efficiency in terms of MFLOPS/W based on HPL runs.

3.4 Intel Xeon-E5

Xeon-E5 is a family of processors that are commonly used in current supercomputers. Although these processors are very popular, the author could not find any papers that take a closer look at the energy efficiency of a single Xeon-E5 processor performing the HPL benchmark. The only values that could be found were the ones provided in the Green500 list. In the November 2015 Green500 list, the "Falcon" supercomputer from the Idaho National Laboratory has the highest energy-efficiency among computers that don't use accelerators. This system has an energy-efficiency of $2262,11$ MFLOPS/W, using Intel Xeon E5-2680v3 CPUs [11] with a clock-frequency of 2.5 GHz. Since the benchmark was done on the whole system, energy-consumption of the interconnect and other parts of the system are part of the measured values. Therefore it can be expected that the energy-efficiency of a single processor is even higher. Intel introduced the Xeon E5-2680v3 in Q3 2014, whereas the before mentioned ARM systems were introduced in 2012 and earlier. Taking a look at the November 2012 Green500 list shows that the by then most energy-efficient supercomputer that uses Intel Xeon-E5 CPUs, an IBM iDataPlex DX360M4 at the Center for Development of Advanced Computing (C-DAC), USA, has an energy-efficiency of $974,69$ MFLOPS/W. It uses Intel Xeon E5-2670 CPUs [10] that were introduced in Q1 2012 and is due to its release date probably more suited for a comparison with the older ARM Cortex-A8/9 processors. Again, it can be expected that the energy-efficiency of a single processor is higher.
Even though the energy-efficiency values of the different systems can hardly be compared due to the reasons mentioned above, some conclusions can be made. First, regarding floating point operations, ARM processors of the ARMv7-A architecture family are less energy-efficient than Intel processors. One reason for this gap in energy-efficiency between ARM CortexA8/9 processors and Intel Xeon are the fewer and smaller floating point registers available on ARM. That the energy-efficiency of ARM processors can be very close or even better than that of Intel Xeon processors depending on the task was shown in [28], where energy-to-solution in *Wh* was used as measure. Secondly, over the years there has been an increase of the energy-efficiency of Intel processors. It can be expected that newer ARM processors will also increase energy-efficiency in the future. Table 1 gives an overview of the achieved energy-efficiency of the different processors.

4 ARM-SoC Servers and HPC

In the previous sections it could be seen that ARM processors are not yet as energy-efficient as Intel processors, but they have potential to be at least as or even more energy-efficient than Intel processors in the future. Nevertheless the overall performance in FLOPS of ARM CPUs is much lower than the performance of Intel CPUs. That means that more nodes need to be switched together in a supercomputer to reach the same performance as a supercomputer running on Intel cores. Therefore it is important to find out, if computers that run on mobile SoCs can scale at least as good as other supercomputer systems or what limits them for the use in supercomputers.
In the last years several products and projects appeared, that make use of ARM processors for server systems and for high-performance computing. Although the use cases for server systems are usually different to that of supercomputers, e.g. database or web hosting, instead of number crunching, server nodes have one thing in common with the nodes used in supercomputers: they usually consist of many CPUs that are switched together and run in parallel. Therefore investigation on the energy-efficiency and performance of server nodes can give

Table 1: Overview of processor specifications and achieved energy-efficiencies.

	Cortex-A8	Cortex-A9			Intel Xeon E5	
CPU/SoC	TI OMAP 3530	TI OMAP 4430	Nova A9500	NVIDIA Tegra 2	E5-2680v3	E5-2670
Architecture	ARMv7				x86_64	
#Cores	1	2	2	2	12	8
Clock	600 MHz	1 GHz	1 GHz	1 GHz	2.5 GHz	2.6 GHz
MFLOPS/W	19.749	91.3	235.0	161.5	2262.11	674.69

a hint on how good supercomputer nodes based on ARM CPUs might perform.

4.1 ARM based Servers

Like other businesses, data-center operators usually want to keep costs low. With that in mind, some CPU manufacturers developed ARM-CPUs especially for the server market. Manufacturers of server hardware developed servers that feature those ARM processors.

Two examples for CPUs developed especially for the server market are the X-Gene CPU from Applied Micro [1] and the ThunderX processor family from Cavium [7], both 64-bit ARMv8 CPUs. Cavium combines up to 48 ARM processors together with PCIe, SATA, memory and network controllers to one single SoC. Examples for servers with ARM processors are the HP ProLiant m400 [9] and Boston's Viridis system [5].

4.1.1 Boston Viridis

The Viridis system from Boston Limited is a server that comes in a 2U rack mountable enclosure targeted at the use in data centers. It contains 48 quad-core Cortex-A9 driven Server-on-Chip modules from Calxeda Inc., each running at a frequency of 1.4 GHz and equipped with 4 GB DDR DRAM. This makes 192 single cores and 192 GB of RAM in one enclosure.

Jarus et al. [23] compared the Boston Viridis to three different configurations of the "CoolEmAll RECS platform"[4] of the Poznan Supercomputing and Networking Center, Poland, and to a Bull BullX S6030. Table 2 lists the different processors that were used in the tested systems. Two of the RECS systems feature processors that are low-power processors (Intel Atom, AMD Fusion), that were originally designed for energy efficiency.

Table 2: Different processors and amount of cores of the systems compared by Jarus et al. [23].

Name	Processor	Cores
RECS	Intel Core i7-3615QE @ 2.3 GHz	18 × 8
RECS	Intel Atom N2600 @ 1.6 GHz	18 × 2
RECS	AMD Fusion G-T40N @ 1 GHz	18 × 2
BullX	Intel Xeon E7-4850 @ 2 GHz	16 × 10
Viridis	ARM Cortex-A9 @ 1.4 GHz	48 × 4

On all listed systems several benchmarks were performed. One of the used benchmarks, besides some single-threaded applications, was the HPL benchmark over the whole system. The approach was, to first find out the best parameters for the HPL run on the different systems and then run HPL with this parameters 100 times. For the BullX system the optimal parameters were already provided by the vendor. Except for the BullX system, where the Intel compiler was used, all HPL binaries were built with the GNU compiler suite.

The results of the HPL runs are listed in table 3. As can be seen, the BullX system with its Intel Xeon E7 CPUs is the best performing system, with 1072 GFLOPS. This is twice as much than the RECS system equipped with Intel Core i7 CPUs, albeit the number of cores differs only by 16 and the RECS i7 runs at a higher clock-rate. One reason for this performance gap of course might be the use of the Intel compiler and the optimized HPL parameters on the BullX. The most energy-

[4]RECS = Resource Efficient Computing System

Table 3: Results of the HPL full platform runs by Jarus et al. [23].

Name	GFLOPS	MFLOPS/W
RECS i7	452.1	561.163
RECS atom	12.76	54.658
RECS AMD	18.85	65.603
BullX	1072	321.74
Viridis	34.39	572.34

efficient system is the Boston Viridis, achieving 572.34 MFLOPS/W. However, it is leading with only a couple of MFLOPS/W in front of the RECS i7. But, since the paper does not provide any information on the used compiler flags or on the usage of SIMD extensions, it is not quite clear if the compiler makes use of all the features available on a system. Therefore it might be that for example no SIMD extensions were used on some of the systems, so performance and energy efficiency are actually even better.
An energy-efficiency of 572.34 MFLOPS/W is still far away from the 7031.58 MFLOPS/W of the "Shoubu" and even more from the anticipated 5 GFLOPS/W. Looking only at the energy-efficiency, not considering the overall performance of the system, the Viridis system would rank at position 288 of the November 2015 Green 500 List. Taking only those systems into account that do not make use of special accelerator hardware, the Viridis would rank at position 206.

4.2 ARM based HPC systems

Currently there are only a few projects ongoing, that aim at building HPC hardware with ARM CPUs. For example the company EUROTECH S.p.A. announced that it is building a HPC architecture based on 64-bit ARM CPUs and NVIDIA GPUs [8]. Lenovo and The Hartree Centre of UK's Science and Technology Facilities Council are jointly working on a research project for developing ARM-based HPC systems [12]. Unfortunately there is only so much information on these projects and no research has been published regarding those systems.

The most advanced project in this field is the Mont-Blanc Project of the European Union, coordinated and hosted by the Barcelona Supercomputing Center, Spain. So far, this project actively runs a prototyped supercomputer based on ARM CPUs.

4.2.1 Mont-Blanc Project

Start of the Mont-Blanc project was in 2011. On the way to develop a large scale supercomputer, a first HPC node prototype had been developed, called Tibidabo [29]. In a later project phase a productive prototype of Mont-Blanc had been deployed.

The Tibidabo system is a cluster of 128 nodes, each featuring a NVIDIA Tegra 2 SoC with a dual-core ARM Cortex-A9 CPU @ 1GHz. The whole system achieves during the execution of HPL an energy-efficiency of 120 MFLOPS/W. This is about 74% of the energy-efficiency of a single node.

In mid 2015, the first Mont-Blanc prototype was successfully deployed [13]. The production partition of the project consists of 7 chassis with 945 compute nodes—15 nodes dedicated to login and 930 dedicated to computation [20]. Every single node consists of a so-called Samsung Daughter Board (SDB), hosting a Samsung Exynos 5250 SoC with the following properties: dual-core Cortex-A15 CPUs @ 1.7 GHz, an ARM Mali T604 GPU with four shader cores (supporting OpenCL 1.1 Full Profile), 4 GB LPDDR3 RAM and 1 Gbit Ethernet. This sums up to 1890 CPU cores and an additional 3780 GPU cores. The SMBs are connected to Ethernet Mother Boards (EMB). Each EMB holds 15 SMBs and an Ethernet switch that connects the 15 SMBs with each other via 1 Gbit/s and via two 10 Gbit/s ports with other EMBs. Unfortunately at the time of writing this paper, there were no publications available that investigated the performance and energy-efficiency of the Mont-Blanc prototype.

5 Conclusions

The current state of research shows that the ARMv7-A based CPUs are less energy-efficient than other processors currently used in HPC systems. This is mainly related to the different design goals of the processor architectures, i.e. while ARM processors were initially developed for the low power mobile and embedded market, the other processors were developed for server and high-performance applications. But with the mobile and embedded market in mind, there was no reason for ARM to focus on operations usually necessary in HPC, like floating point operations. Thus the capabilities of those processors are very limited in this particular case. With the newer ARMv8-A based processors ARM is closing the performance gap to Intel processors. The ARMv8-A architecture brings more power and a richer feature set into ARM processors, which makes them more competitive, but still at a lower energy consumption.

ARM processors already make their way into HPC. With the European Mont-Blanc project, ARM processors have taken a huge step to enter the HPC market. Also for the server market ARM seems to be a promising alternative to Intel processors. The fact that first ARM-based supercomputers are built shows that a cluster of ARM processors or SoCs can potentially scale up to large systems and ARM is indeed suitable for HPC.

What hasn't been a part of the research so far is how the ARM processors perform when they are used together with special accelerators, like GPGPUs or Intel's Xeon Phis. Another interesting point is the fact that some of the newer on-board graphic chips of the ARM SoCs are also capable of general purpose processing, by enabling access through the OpenCL framework. So far, the only ones who did research on the programmability of ARM processors using OpenCL were Richie et al. [31]. Although they showed the general programmability through OpenCL, their approach did not go beyond the usage of the CPUs of a Calxeda ARM Server, which has no GPGPUs.

References

[1] Applied Micro Circuits Corporation, X-Gene®. https://www.apm.com/products/data-center/x-gene-family/x-gene/. Accessed: 2015-12-05.

[2] ARM, ARM® Cortex®-A15 MPCore™ Processor, Technical Reference Manual. http://infocenter.arm.com/help/topic/com.arm.doc.ddi0438i/DDI0438I_cortex_a15_r4p0_trm.pdf. Revision: r4p0.

[3] ARM, Cortex™-A8, Technical Reference Manual. http://infocenter.arm.com/help/topic/com.arm.doc.ddi0344k/DDI0344K_cortex_a8_r3p2_trm.pdf. Revision: r3p2.

[4] ARM, Cortex™-A9 Floating-Point Unit, Technical Reference Manual. http://infocenter.arm.com/help/topic/com.arm.doc.ddi0408i/DDI0408I_cortex_a9_fpu_r4p1_trm.pdf. Revision: r4p1.

[5] Boston Viridis Data Sheet. http://download.boston.co.uk/downloads/9/3/2/932c4ecb-692a-47a9-937d-a94bd0f3df1b/viridis.pdf. Accessed: 2015-12-07.

[6] CALAO Systems SKY-S9500-ULP-CXX (aka SNOWBALL). http://www.calao-systems.com/repository/pub/SINGLE%20BOARD%20COMPUTERS/SKY-S9500-ULP-XXX/DOCS/SKY-S9500-ULP-CXX-FLY-V1.3.pdf. Accessed: 2015-12-03.

[7] Cavium, ThunderX™. http://www.cavium.com/Table.html#thunderx. Accessed: 2015-12-18.

[8] Eurotech Develops New High-Performance Computing Architecture Combining Applied Micro 64-bit ARM CPUs and NVIDIA GPU Accelerators. http://www.eurotech.com/en/press+room/news/?674. Accessed: 2015-12-04.

[9] Hewlett Packard Enterprise Development LP, HP ProLiant m400.

http://www8.hp.com/us/en/products/proliant-servers/product-detail.html?oid=7398907#!tab=specs. Accessed: 2015-12-18.

[10] Intel®Xeon®Processor E5-2670 (20M Cache, 2.60 GHz, 8.00 GT/s Intel®QPI), specifications. http://ark.intel.com/products/64595/Intel-Xeon-Processor-E5-2670-\20M-Cache-2_60-GHz-8_00-GTs-Intel-QPI. Accessed: 2015-12-11.

[11] Intel®Xeon®Processor E5-2680 v3 (30M Cache, 2.50 GHz), Specifications. http://ark.intel.com/products/81908/Intel-Xeon-Processor-E5-2680-v3-\30M-Cache-2_50-GHz. Accessed: 2015-12-11.

[12] Lenovo and U.K-based Hartree Centre to Advance Energy-Efficient Computing. http://news.lenovo.com/news+releases/lenovo-and-hartree-centre-advance-\energy-efficient-computing.htm#rel. Accessed: 2015-12-04.

[13] Mont-Blanc Consortium Partners, D7.8 Prototype System Deployed. https://www.montblanc-project.eu/sites/default/files/MB_D7.8_Prototype%20System%20Deployed%20v1.pdf. Accessed: 2015-12-04.

[14] Reduced complexity, increased profitability. http://www.energy.siemens.com/mx/pool/hq/power-generation/renewables/wind-power/platform%20brochures/adaption-d3-onshore-brochure-en.pdf. Accessed: 2015-12-10.

[15] The Graph 500 List. http://www.graph500.org/. Accessed: 2015-12-02.

[16] The Green Graph 500 List. http://green.graph500.org/. Accessed: 2015-12-02.

[17] The Green500 List. http://www.green500.org/. Accessed: 2015-12-10.

[18] The TOP500 list, November 2010. http://top500.org/lists/2010/11/. Accessed: 2016-01-31.

[19] The TOP500 list, November 2015. http://top500.org/lists/2015/11/. Accessed: 2016-01-31.

[20] Wiki-Pages of the Mont-Blanc Project, Mont-Blanc Prototype. https://wiki.hca.bsc.es/wiki/MontBlanc/?animal=MontBlanc. Accessed: 2015-12-04.

[21] *High Performance Computing in Science and Engineering*, Garching/Munich, 2014. Siegfried Wagner, Arndt Bode, Helmut Satzger, Matthias Brehm.

[22] E. Blem, J. Menon, and K. Sankaralingam. Power struggles: Revisiting the risc vs. cisc debate on contemporary arm and x86 architectures. In *High Performance Computer Architecture (HPCA2013), 2013 IEEE 19th International Symposium on*, pages 1–12, Feb 2013.

[23] Mateusz Jarus, Sbastien Varrette, Ariel Oleksiak, and Pascal Bouvry. Performance evaluation and energy efficiency of high-density hpc platforms based on intel, amd and arm processors. In Jean-Marc Pierson, Georges Da Costa, and Lars Dittmann, editors, *Energy Efficiency in Large Scale Distributed Systems*, volume 8046 of *Lecture Notes in Computer Science*, pages 182–200. Springer Berlin Heidelberg, 2013.

[24] Peter Kogge, Keren Bergman, Shekhar Borkar, Dan Campbell, William Carlson, William Dally, Monty Denneau, Paul Franzon, William Harrod, Jon Hiller, Sherman Karp, Stephen Keckler, Dean Klein, Robert Lucas, Mark Richards, Al Scarpelli, Steven Scott, Allan Snavely, Thomas Sterling, R. Stanley Williams, and Katherine Yelick. Exascale computing study: Technology challenges in achieving exascale systems, 2008.

[25] E Padoin, D de Olivera, Pedro Velho, and P Navaux. Evaluating energy efficiency and instantaneous power on arm plafforms,. In *Proc. of 10th Workshop on Parallel and Distributed Processing*, 2012.

[26] E Padoin, D de Olivera, Pedro Velho, P Navaux, Brice Videau, Augustin Degomme, and Jean-Francois Mehaut. Scalability and energy efficiency of hpc cluster with arm mpsoc. In *Proc. of 11th Workshop on Parallel and Distributed Processing*, 2013.

[27] E.L. Padoin, D.A.G. de Oliveira, P. Velho, and P.O.A. Navaux. Evaluating performance and energy on arm-based clusters for high performance computing. In *Parallel Processing Workshops (ICPPW), 2012 41st International Conference on*, pages 165–172, Sept 2012.

[28] E.L. Padoin, D.A.G. de Oliveira, P. Velho, and P.O.A. Navaux. Time-to-solution and energy-to-solution: A comparison between arm and xeon. In *Applications for Multi-Core Architectures (WAMCA), 2012 Third Workshop on*, pages 48–53, Oct 2012.

[29] Nikola Rajovic, Alejandro Rico, Nikola Puzovic, Chris Adeniyi-Jones, and Alex Ramirez. Tibidabo: Making the case for an arm-based hpc system. *Future Generation Computer Systems*, 36:322 – 334, 2014. Special Section: Intelligent Big Data ProcessingSpecial Section: Behavior Data Security Issues in Network Information PropagationSpecial Section: Energy-efficiency in Large Distributed Computing ArchitecturesSpecial Section: eScience Infrastructure and Applications.

[30] Nikola Rajovic, Alejandro Rico, James Vipond, Isaac Gelado, Nikola Puzovic, and Alex Ramirez. Experiences with mobile processors for energy efficient hpc. In *Proceedings of the Conference on Design, Automation and Test in Europe*, DATE '13, pages 464–468, San Jose, CA, USA, 2013. EDA Consortium.

[31] David Richie, James Ross, Jordan Ruloff, Song Park, Lori Pollock, and Dale Shires. Investigation of parallel programmability and performance of a calxeda arm server using opencl. In Dieter an Mey, Michael Alexander, Paolo Bientinesi, Mario Cannataro, Carsten Clauss, Alexandru Costan, Gabor Kecskemeti, Christine Morin, Laura Ricci, Julio Sahuquillo, Martin Schulz, Vittorio Scarano, StephenL. Scott, and Josef Weidendorfer, editors, *Euro-Par 2013: Parallel Processing Workshops*, volume 8374 of *Lecture Notes in Computer Science*, pages 865–874. Springer Berlin Heidelberg, 2014.